

How Hidden are Hidden Processes?

A Primer on Crypticity and Entropy Convergence

John R. Mahoney,^{1,*} Christopher J. Ellison,^{2,†} Ryan G. James,^{2,‡} and James P. Crutchfield^{2,3,§}¹*Physics Department, University of California at Merced,
5200 North Lake Road Merced, CA 95343*²*Complexity Sciences Center
Physics Department, University of California at Davis,
One Shields Avenue, Davis, CA 95616*³*Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501*

(Dated: August 9, 2011)

We investigate a stationary process's crypticity—a measure of the difference between its hidden state information and its observed information—using the causal states of computational mechanics. Here, we motivate crypticity and cryptic order as physically meaningful quantities that monitor how hidden a hidden process is. This is done by recasting previous results on the convergence of block entropy and block-state entropy in a geometric setting, one that is more intuitive and that leads to a number of new results. For example, we connect crypticity to how an observer synchronizes to a process. We show that the block-causal-state entropy is a convex function of block length. We give a complete analysis of spin chains. We present a classification scheme that surveys stationary processes in terms of their possible cryptic and Markov orders. We illustrate related entropy convergence behaviors using a new form of foliated information diagram. Finally, along the way, we provide a variety of interpretations of crypticity and cryptic order to establish their naturalness and pervasiveness. Hopefully, these will inspire new applications in spatially extended and network dynamical systems.

Keywords: crypticity, stored information, statistical complexity, excess entropy, information diagram, synchronization, irreversibility

PACS numbers: 02.50.-r 89.70.+c 05.45.Tp 02.50.Ey 02.50.Ga

		B. State-block entropy	8
		C. Block-state entropy	9
CONTENTS		VI. Crypticity in Spin Chains	10
I. Introduction	2	VII. Geometric Constraints	12
II. Definitions	3	VIII. The Cryptic Markovian Zoo	12
III. Crypticity: From state paths to synchronization	3	IX. Information Diagrams for Stationary Processes	13
A. Global versus Local	5	X. Conclusion	15
B. Mazes and Stacks	5	Acknowledgments	15
C. Transient versus Relaxed	5	A. Why Crypticity?	15
D. Naive versus Informed	6	B. Equivalence of Forward and Reverse Restricted State-Paths	16
E. Statistical Complexity versus Crypticity	6	C. Crypticity and Co-unifilarity	17
IV. Crypticity through Information Theory	6	References	17
A. Crypticity	6		
B. Cryptic Order	7		
V. Crypticity and Entropy Convergence	8		
A. Block Entropy	8		

* jmahoney3@ucmerced.edu

† cellison@cse.ucdavis.edu

‡ rgjames@ucdavis.edu

§ chaos@cse.ucdavis.edu

A black box is a metaphor for ignorance: One cannot see inside, but the presumption is that something, unknown in whole or in part, is there to be discovered. Moreover, the conceit is that the impoverished outputs from the box do contain something partially informative. Physically, ignorance comes in the act of measurement—measurements that are generically incomplete, inaccurate, and infrequent. Since measurements dictate that one can have only a partial view, it goes without saying that these distortions make discovery both difficult and one of the key challenges to scientific methodology. Measurement necessarily leads to our viewing the world as being hidden from us. Of course, the world is not completely hidden. If it were, then there would be neither gain nor motivation to probing measurements to build models. Scientific theory building and its experimental verification operate, then, in the framework of hidden processes—processes from which we have observations from which, in turn, we attempt to understand the hidden mechanisms. At least philosophically, this setting is not even remotely new. The circumstance is that addressed by Plato’s metaphor of our knowledge of the world deriving from the data of shadows on a cave wall.

Fortunately, we are far beyond metaphors these days. Hidden processes pose a quantitative question: How hidden are they? Here, we show how to quantitatively measure just this: How much internal information is hidden by measuring a process? Of course, this assumes, as in the black box metaphor, that there is something to be discovered. The tool we use to ground the intentional stance of discovering the internal mechanisms—to say *what* is hidden—is computational mechanics. Computational mechanics is a theory of what patterns are and how to measure a hidden process’s degree of structure and organization. Computational mechanics has a long history, though, going back to the original challenges of nonlinear modeling posed in the 1970s that led to the concept of reconstructing “geometry from a time series”. The explorations here can be seen in this light, with one important difference: Computational mechanics shows that measurements of a hidden process tell how the process’s internal organization should be represented. Building on this, we develop a quantitative theory of how hidden processes are.

I. INTRODUCTION

Many scientific domains face the confounding problems of defining and measuring information processing in dynamical systems. These range from technology to fundamental science and, even, epistemology of science [1]:

1. *The 2020 Digital Roadblock*: The end of Moore’s scaling laws for microelectronics [2–4].
2. *The Central Dogma of Neurobiology*: How are the intricate physical, biochemical, and biological components structured and coordinated to support natural, intrinsic neural computing?
3. *Physical Intelligence*: Does intelligence require biology, though? Or can there be alternative nonbiological substrates which support system behaviors that are to some degree “smart”.
4. *Structure versus Function*: Intelligence aside, how do we define and detect spontaneous organization, in the first place? How do these emergent patterns take on and support functionality?

Many have worked to quantify various aspects of information dynamics; cf. Ref. [5]. One often finds references to information storage, transfer, and processing. Sophisticated measures are devised to characterize these quantities in multidimensional settings, including networks and adaptive systems.

Here, we investigate foundational questions that bear on all these domains, using methods with very few modeling and representation requirements attached that, nonetheless, allow a good deal of progress. In quantifying information processing in stochastic dynamical systems, two measures have repeatedly appeared and been successfully applied: the past-future mutual information of observations (excess entropy) \mathbf{E} [6, and references therein] and the internal stored information (statistical complexity) C_μ [7]. Curiously, the difference between these measures—the crypticity χ [8]—has only recently received attention. To our knowledge, the first attempt to understand χ directly was in Ref. [9]. The following provides additional perspective and clarity to the results contained there and in the related works of Refs. [8, 10, 11]. In particular, we add to the body of knowledge surrounding crypticity and cryptic order, develop a further classification of the space of processes, and introduce several alternative ways to visualize these concepts. An appendix demonstrates that crypticity captures a notable and unique property, when compared to alternative information measures. The goal is to provide a more intuitive and geometric toolbox for posing and answering the increasing range and increasingly

more complex research challenges surrounding information processing in nature and technology.

II. DEFINITIONS

We denote contiguous groups of random variables X_i using $X_{n:m+1} = X_n \dots X_m$. A semi-infinite group is denoted either $X_{n:} = X_n X_{n+1} \dots$ or $X_{:n} = \dots X_{n-2} X_{n-1}$. We refer to these as the *future* and the *past*, respectively. Consistent with this, the bi-infinite chain of random variables is denoted $X_{:}$. A *process* is specified by the distribution $\Pr(X_{:})$. Throughout the following, we assume we are given a stationary process.

Please refer to Refs. [9, 12] for supplementary definitions of presentations, causal states, ϵ -machines, unifiarity, co-unifiarity, Shannon block information, information diagrams, and the like. The following assumes familiarity with these concepts and the results and techniques there. However, our development calls for a few reminders.

There are two notions of memory central to characterizing stochastic processes. These are the excess entropy \mathbf{E} (sometimes called the predictive information) and the statistical complexity C_μ . The excess entropy is a measure of correlation between the past and future: the degree to which one can remove uncertainty in the future given knowledge of the past. (This is illustrated as the green information atom at the intersection the past and future in the information diagram of Fig. 1.) The statistical complexity is a quantity that arises in the context of modeling rather than prediction. Specifically, C_μ is the amount of information required for an observer to synchronize a stochastic process. In the setting of finite-state hidden Markov models, it is the information stored in the process's causal states.

Then, we have the crypticity:

Definition 1. A process's crypticity χ is defined as:

$$\chi = H[S_0 | X_{0:}] ,$$

where S_t is a process's causal state at time t .

Clearly, the definition relies on having a process's ϵ -machine presentation; the states used are causal states. Other presentations, whose alternative states we denote \mathcal{R} , suggest an analogous, but more general definition of crypticity; cf. Ref. [12].

To give us something to temporarily hang our hat on, it turns out that the crypticity is simply how much stored information is hidden from observations. That is, it is the difference between the internal stored information (C_μ) and the apparent past-future mutual information (\mathbf{E}). This is directly illustrated in Fig. 1.




FIG. 1. Crypticity χ is represented by the red (dark) crescent shape in this ϵ -machine I-diagram. The excess entropy \mathbf{E} , by the (green) overlap of the past information $H[X_{:0}]$ and future information $H[X_{0:}]$. The statistical complexity C_μ is the information in the internal causal states \mathcal{S} and comprises both χ and \mathbf{E} . For a review of information measures and diagrams refer to the citations given in the text or quickly read the first portions of Sec. IX.

We are also interested in the range required to “learn” the crypticity. This is the *cryptic order*.

Definition 2. A process's cryptic order k is defined as:

$$k = \min\{L \in \mathbb{Z}^+ : H[S_L | X_{0:}] = 0\} .$$

These definitions do not easily admit an intuitive interpretation. Their connection to hidden stored information is not immediately clear, for example. They mask the importance and centrality of the crypticity property. Given this, we devote some effort in the following to motivate them and to give several supplementary interpretations. As a start, Fig. 1 gives a graphical definition of crypticity using the ϵ -machine information diagram of Ref. [8]. It is the red crescent highlighted there, which is the state information $C_\mu = H[\mathcal{S}]$ minus that information derivable from the future $H[X_{:0}] \equiv H[X_{0:}]$. This begins to explain crypticity as a measure of a process's hidden-ness. We'll return to this, but first let's consider several other alternatives.

III. CRYPTICITY: FROM STATE PATHS TO SYNCHRONIZATION

Crypticity and, in particular, cryptic order have straightforward interpretations when one considers the internal state-paths taken as an observer synchronizes to a process [13]. In this, cryptic order is seen to be analogous to, and potentially simpler than, a process's Markov order. While both the Markov and cryptic orders derive from a notion of synchronization, the cryp-

tic order depends on a subset of the paths realized during synchronization. We illustrate this via an example: The (R, k) -Golden Mean Process—a generalization of the Golden Mean Process with tunable Markov order R and tunable cryptic order k . In particular, we examine the $(3, 2)$ -Golden Mean Process shown in Fig. 2.




FIG. 2. The $(3, 2)$ -Golden Mean ϵ -machine: Markov order 3 and cryptic order 2.

It is straightforward to verify that the only words of length 3 generated by this process are $\{000, 001, 011, 100, 110, 111\}$. Since the process is Markov order 3 (by construction) we know that each of these words is a synchronizing word [14]. Some words lead to synchronization in fewer than three steps, though. For instance, 011 yields synchronization to state E after just the first two symbols 01.

In Fig. 3, we display the internal-state paths taken by each possible initial state under evolution governed by the six synchronizing words. Let's take a moment to describe these illustrations carefully. Before reading any word, there is maximum uncertainty in the internal state. We represent this using a circle for each of the five causal states of the ϵ -machine. Each of these states is led to a next state by following the first symbol seen [15]. For word 001, the first symbol is 0, and A , for instance, is led to B . Notice that E is not led to any state. This is because E has no outgoing transition with symbol 0. The path from E , therefore, ends and is not considered further. The termination of paths is one of the important features of synchronization to note.

Looking at the synchronizing word 100, we see that the transition on the first symbol 1 takes both states A and E to the same state A . Since we use unifilar presentations (ϵ -machines), this merging can never be undone. Path merging is yet another important feature.

Both the termination and merging of paths are relevant to synchronization, but have different roles in the determination of the Markov and cryptic orders.

Although we already know the Markov order of this process, we can read it from Fig. 3 by looking at the lengths for each word where only one path remains. These lengths $\{3, 3, 2, 2, 2, 2\}$ are marked with orange dia-

monds. The maximum value of this length is the Markov order (3, in this example).




FIG. 3. Synchronization paths for $(3, 2)$ -Golden Mean ϵ -machine: Each synchronizing word induces a set of state-paths; some of which terminate, some of which merge.

In the next illustration, Fig. 4, we keep only those paths that do not terminate early. In this way, we remove paths that generally are quite long, but that terminate before having the chance to merge with the final synchronizing paths. We similarly mark, with green triangles, the lengths where these reduced paths have ultimately merged. Note that restricting paths can only preserve or decrease each length. Finally, in analogy to the Markov order, the maximum of these lengths $\{0, 0, 0, 1, 2, 2\}$ is the cryptic order (2 in this example).




FIG. 4. The paths that are not terminated before the Markov order are highlighted in red. These are the paths relevant for the cryptic order. For each word the contribution to Markov order is still indicated by an orange diamond, whereas the contribution toward cryptic order is indicated by a green triangle.

This demonstrates how crypticity relates to paths and path merging. It is a small step then to ask for a direct connection to co-unifilarity [9]: $H[S_0|X_0S_1] = 0$. In fact, there are three primary equivalent statements about a process: (i) its ϵ -machine being co-unifilar, (ii) its $\chi = 0$, and (iii) its cryptic order $k = 0$. (Appendix C presents a proof of this equivalence in terms of entropy growth functions and includes the connection to cryptic order as well.)

This exposes the elementary nature of the cryptic order as a property of synchronizing paths. Appendix B goes further to show that state-paths traced similarly, but in the reverse time direction, are the same as those singled-out in the forward direction, as just done. The remainder of this section offers different perspectives on crypticity, some of which are less strict, but provide intuition and suggest its broad applicability.

A. Global versus Local

Imagine a synchronization task involving a group of agents. The agents begin in different locations (states) and move to next locations based on the synchronization input they receive from a common controller. The goal is to provide a uniform input that causes (a subset of) the agents to arrive at the same location. This is reminiscent of a road-coloring problem. In many road-coloring contexts, only uniform-degree graph structures are investigated, largely due to theoretical tractability. However, real-world graphs are rarely uniform degree. This means that some agents may receive instructions that they cannot carry out. These agents quit, and their paths are terminated. Assuming that the instructions are synchronizing for some subset of the agents (the instruction is a synchronizing word), the synchronization task will end with this subset of agents at the desired destination.

There are two ways in which we may view this process. One is global, and corresponds to the Markov order, while the other is local and corresponds to the cryptic order.

If we monitor the entire collection of agents from a bird’s eye view evolving under the synchronization input, we observe paths terminating and merging. Our global notion of synchronization is the point at which each path is either terminated or merged with every other valid path. This is clearly coincident with the description of Markov order previously described.

Alternatively, we monitor the collective by querying the agents after the task is complete. The unsuccessful agents, whose paths were terminated, never arriving at the destination, cannot be queried. From this viewpoint, synchronization takes place *relative* to the group of agents that were not terminated. As locally interacting entities, they know the latest time at which an agent merged with their group—the group which ultimately synchronized. Even after this event, there may be other agents still operating that will inevitably be terminated at some later time. This means that from the local (agent) perspective, synchronization may happen earlier than from the global (controller) perspective.

We claim, based on this setting, that the cryptic order has a straightforward and physically relevant basis in the context of synchronization. Upcoming discussions, some more technical, will emphasize this point further, as well as demonstrate new results.

B. Mazes and Stacks

The Markov versus cryptic order distinction is relevant to any maze-solving algorithm [16]. Imagining the solution of a maze as a sequence of moves—left, right, or

straight—we may write down a list of potential solutions (which must contain all actual solutions) by listing all 3^{N^2} sequences [17]. A brute-force algorithm tries all of these paths. Since we are interested in worst-case scenarios, many of the details (e.g., depth- versus breadth-first search) are not relevant. What *is* relevant is the object that the algorithm must maintain in memory or that it ultimately returns to the user.

An algorithm might try out each potential solution, feeding in each move sequentially and testing for either maze completion or termination (walking into a wall or a previously visited location) at each step. The end of each solution is marked with a length. When all solutions have been tried, this set of solutions and lengths is returned. While this is not a stationary stochastic process, we may think of the longest of these lengths as being similar to the Markov order. The speed and memory use of this algorithm are obviously improved by using a tree structure, but this does not affect the result we are interested in.

If we were only interested in paths which end in maze completion, an even more memory-conscious algorithm would realize that dead-ends in the tree could be removed. One accomplishes this with a stack memory for the active-path tree branch. Reaching a nonsolving termination, the algorithm pops the end states until returning to the most recent unexplored option. This process continues recursively until the tree has been filled out. The relevant lengths are now the lengths of the maze-completing paths (all root-to-leaf paths), the longest of which is an analog of the cryptic order.

C. Transient versus Relaxed

Rather than using the global versus local distinction, we can think in terms of a dynamical view of synchronization. We might imagine a collection of ants attempting to create paths from a resource-rich region to their nest; or a watershed in the process of forming the transport network from collection regions to the main body of water. Until these networks develop, it is not clear which will become the important paths.

A log not worth climbing over causes ants to make the effort less often, thereby dropping less pheromone, leading fewer ants to attempt this path, until finally it is empty. Similarly, slow water deposits more sediment and fills underused channels. As these networks evolve from an initial transitory state to relaxed state, the types of paths within the network and their synchronization properties change. In particular, while the early-time synchronization depends on the terminating paths, the later-time synchronization will not. In this dynamical picture we

see that a property akin to cryptic order emerges as the system evolves.

D. Naive versus Informed

It is only a small step from this dynamical picture to view these self-reinforcing systems as evolving from naive to informed states. Over time, a system “realizes” which paths are undesirable and quits them. Consider an individual learning to navigate a new city. She will experience a similar network evolution, where the pruning of dead-end paths is an intentional act. This navigation structure also will tend to reflect the cryptic order.

E. Statistical Complexity versus Crypticity

In addition to describing the Markov and cryptic orders via a dynamical picture of synchronization, we can explore the same phenomenon with the associated entropies, a more statistical perspective.

Beginning with the global view, the distribution over the set of all starting points is the state entropy $H[\mathcal{S}_0]$, commonly called the statistical complexity C_μ . By considering the initial state distribution conditioned on the removal of the terminating paths, we are left with only a portion of this entropy, and this is the crypticity χ [18]. As discussed, we might consider this removal a result of memory, relaxation, or prescience.

IV. CRYPTICITY THROUGH INFORMATION THEORY

The discussion above in terms of paths is relatively intuitive. The original conception, however, was not in terms of paths, but rather in terms of information-theoretic quantities. Information identities based on ϵ -machines are beginning to provide a growing set of interpretations; some more subtle, some more direct than others. The following will show that crypticity and cryptic order have diverse implications and also that even elementary information-theoretic quantities form a rich toolset.

A. Crypticity

The ϵ -machine causal presentation pairs up pasts with futures in a way appropriate for prediction. Since pasts can be different but predictively equivalent, this pairing operates on sets of pasts that, in turn, are equivalent to the causal states themselves. Furthermore, a single past

can be followed by a set of futures. This is natural since the processes are stochastic. So, any past or predictively equivalent group of pasts is linked to a distribution of futures. Finally, these future distributions often overlap. As we will now show, crypticity is a measure of this overlap.

Historically, it has taken some time to sort out the similarities and differences between various measures of memory. Eventually, two emerged naturally as key concepts: C_μ , the statistical complexity or information processing “size” of the internal mechanism; and \mathbf{E} , the excess entropy, or the apparent (to an observer) amount of past-future mutual information. It has been recognized for some time [19, 20] that C_μ is an upper bound on \mathbf{E} . The strictness of this inequality and the nature of the relationship between the two, however, was not significantly explored until Ref. [9]. The first simple statement [8] about crypticity in terms of information-theoretic quantities is that it is the quantifiable difference between two predominant measures of information storage: $\chi = C_\mu - \mathbf{E}$.

Taking this view a bit further, since \mathbf{E} is the amount of uncertainty in the future that one can reduce through study of the past, and C_μ is the amount of information necessary to do optimal prediction (using a minimal predictor), their difference is the amount of *modeling overhead*. One may object that a minimal optimal predictor should not require more information than will be made use of. In fact it is known that many processes with large χ have nonunifilar representations that are much smaller [9]. What is not obvious is that this is simply a re-representation of the causal states as mixtures of the new states [11, 21]. In other words, *the overhead is inescapable*. This suggests a useful language with which to discuss stochastic processes—not only do we identify a process with an ϵ -machine, but we analyze the efficiency of these machines in terms of required resources.

For the following, we briefly invoke the use of the reverse ϵ -machine, the causal representation of a process when scanned in reverse, to extend our view of crypticity. (For details on reverse causal states, see Refs. [8, 11].) Recall that forward causal states are built for prediction and, similarly, reverse causal states are built for retrodiction. We say they are “built” for these purposes in the sense that they are minimal and optimal, two desirable design goals.

Given this, it is somewhat surprising to see that forward causal states are better at retrodiction than reverse causal states. The information diagram in Fig. 5 illustrates this. We will now show that the degree to which this is true is precisely the forward process’s crypticity. Here, we write this difference in retrodictive uncertainty




FIG. 5. Crypticity as the degree to which forward causal states are better retrodictors than reverse causal states.

as follows:

$$H[X_{-L:0}|\mathcal{S}_0^-] - H[X_{-L:0}|\mathcal{S}_0^+] \geq 0.$$

Then, this difference converges to χ :

$$\chi = \lim_{L \rightarrow \infty} (H[X_{-L:0}|\mathcal{S}_0^-] - H[X_{-L:0}|\mathcal{S}_0^+]).$$

We might wonder why the reverse causal states were not built to be better at their job. This is explained by the fact that the information input to the above constructs is not equivalent. The forward causal states are built from the past, while the reverse causal states are built from the future. It is no surprise, then, that the forward states can offer information about the pasts from which they were built. It is more interesting to consider why they do not maintain all of this information. This is because the forward states were designed for predicting a stochastic process, a goal for which maintaining information about the past offers diminishing returns.

Rather than comparing the function of two objects (forward and reverse causal states), we can compare two functions of the same object. In this light, the crypticity is the degree to which forward causal states are better at retrodiction than they are at prediction. More precisely, we have:

$$\begin{aligned} H[X_{0:L}|\mathcal{S}_0] - H[X_{0:L}|\mathcal{S}_L] \\ &= H[\mathcal{S}_0|X_{0:L}] - H[X_{0:L}|\mathcal{S}_L] \\ &= H[\mathcal{S}_0|X_{0:L}\mathcal{S}_L] - H[\mathcal{S}_L|\mathcal{S}_0X_{0:L}] \\ &= H[\mathcal{S}_0|X_{0:L}\mathcal{S}_L] \\ &\geq 0. \end{aligned}$$

The first step follows from stationarity, the second appeals to an informational identity, and the next to unifiarity of the ϵ -machine. Similarly, this difference con-

verges to χ :

$$\lim_{L \rightarrow \infty} H[\mathcal{S}_0|X_{0:L}\mathcal{S}_L] = \lim_{L \rightarrow \infty} H[\mathcal{S}_0|X_{0:L}] = \chi$$

Thus, crypticity is the amount of information that, although necessary for current prediction, must be erased at some future time.

B. Cryptic Order

Many of these statements about uncertainty can be rephrased in terms of length scales. The length scale associated with the crypticity is the cryptic order: the distance we must look into the past to discover the modeling overhead. Following our discussion of forward and reverse states, we can interpret cryptic order as the length at which the difference converges to χ :

$$k = \min\{L : H[X_{-L:0}|\mathcal{S}_0^-] - H[X_{-L:0}|\mathcal{S}_0^+] = \chi\}.$$

Stated differently, it is the length at which all advantage of a forward state over a reverse state as a retrodictor is lost. In other words:

$$k = \min\{L : H[X_0|X_{1:L+1}\mathcal{S}_{L+1}^+] = H[X_0|X_{1:L+1}\mathcal{S}_{L+1}^-]\}.$$

Equivalently cryptic order is the length at which a forward state's uncertainty in prediction and retrodiction equalize. More colloquially, it is the range beyond which a forward state is equally good at prediction and retrodiction, or:

$$k = \min\{L : H[X_L|\mathcal{S}_0X_{0:L}] = H[X_0|X_{1:L+1}\mathcal{S}_{L+1}]\}.$$

As Sec. III suggested, the cryptic order k is closely analogous to the Markov order R . Here, we state the parallel formally:

$$\begin{aligned} R &= \min\{L : H[\mathcal{S}_L|X_{0:L}] = 0\} \\ k &= \min\{L : H[\mathcal{S}_L|X_{0:L}, X_L] = 0\}. \end{aligned}$$

Appendix A argues for the uniqueness of this parallel.

Cryptic order is the largest noninferable state sequence. Given an infinite string of measurements $\dots x_{-2}x_{-1}x_0$, one eventually synchronizes to a particular causal state [22], for any finite-state ϵ -machine. The same symbol sequence can then be used to retrodict states beginning at the point of synchronization. All but the earliest k states can be definitively retrodicted regardless of which observed sequence (and resulting predictive state) occurs.

V. CRYPTICITY AND ENTROPY CONVERGENCE

It has become increasingly clear that entropy functions are useful characterizations of processes. Since a process is a bi-infinite collection of random variables [23], it typically is not useful to calculate the entropy of the entire collection. The alternative strategy is to analyze the entropy of increasingly large finite portions. The scaling, then, captures the system's bulk properties in the large-size (thermodynamic) limit, as well as how those properties emerge from the individual components.

These functions capture much of the behavior that we are interested in here. The block entropy $H[X_{0:L}]$ was used to great effect in Ref. [6] to understand the way perceived randomness may be reformulated as structure, when longer correlations are considered. More recently, Ref. [12] used extended functions—the block-state entropy $H[X_{0:L}\mathcal{R}_L]$ and the state-block entropy $H[\mathcal{R}_0X_{0:L}]$ —to explore the relationship between alternate presentations of a process and the information theoretic measures of memory in a presentation.

We will borrow these two new entropy functions and turn them back on the canonical set of presentations, ϵ -machines, to expose the workings of crypticity. The result is a graphical approach that offers a more intuitive understanding of the results originally developed in Ref. [11]. Using this, we sharpen several theorems, discover new bounds, and pose additional challenges.

A. Block Entropy

The *block entropy* $H[X_{0:L}]$ is the joint Shannon entropy of finite sequences. As it is treated rather thoroughly in Ref. [6], we simply recall several of its features.

First, recall that $X_{0,0}$ represents the random variable for a null observation and, since there is just one way to do this, $H[X_{0,0}] = 0$. As L increases, the block entropy curve is a nondecreasing, concave function that limits to the linear asymptote $\mathbf{E} + h_\mu L$, where \mathbf{E} is the excess entropy and h_μ is the process entropy rate.

Given a block entropy curve, Markov processes are easily identified since the curve reaches its linear asymptote at finite block length. That is, the Markov order R satisfies:

$$R \equiv \min \{L : H[X_{0:L}] = \mathbf{E} + h_\mu L\} .$$

Before reaching the Markov order, one has not discovered all process statistics and, so, new symbols appear more surprising than they otherwise would. Mathemati-

cally, this is formulated through a lower bound:

$$H[X_L|X_{0:L}] \geq h_\mu ,$$

for all L . Since the block entropy curve for Markovian processes reaches its asymptote at $L = R$ and since the linear asymptote has slope equal to the entropy rate, we know that Markov processes attain the lower bound whenever $L > R$: $H[X_L|X_{0:L}] = h_\mu$.

Finally, since the block entropy is concave and non-decreasing, it is bounded above by its linear asymptote. This naturally leads to a concave, nondecreasing lower bound estimate for the excess entropy:

$$\mathbf{E}(L) \equiv H[X_{0:L}] - h_\mu L .$$

Thus, $\mathbf{E}(L) \leq \mathbf{E}(L+1) \leq \mathbf{E}$ and $\lim_{L \rightarrow \infty} \mathbf{E}(L) = \mathbf{E}$.

B. State-block entropy

The state-block entropy $H[\mathcal{R}_0X_{0:L}]$ is the joint uncertainty one has in a presentation's internal state \mathcal{R} and the block of symbols immediately following. Its behavior is generally nontrivial, but when restricted to ϵ -machines, its behavior is simple [12]. In that case, it refers to the process's unknown causal state \mathcal{S}_0 and is denoted $H[\mathcal{S}_0X_{0:L}]$.

Its simplicity is a direct consequence of the causal states' efficient encoding of the past. To see this, note that differences in the state-block entropy curve, the rate at which it grows with block length, are constant:

$$\begin{aligned} H[\mathcal{S}_0X_{0:L+1}] - H[\mathcal{S}_0X_{0:L}] &= H[X_L|\mathcal{S}_0X_{0:L}] \\ &= H[X_L|\mathcal{S}_{0:L}X_{0:L}] \\ &= H[X_L|\mathcal{S}_L] \\ &= h_\mu . \end{aligned}$$

Here, we used the unifilarity property of ϵ -machines: $H[\mathcal{S}_{L+1}|\mathcal{S}_L, X_L] = 0$. So, given the causal state \mathcal{S}_0 , the block $X_{0:L}$ of symbols immediately following it determines each causal state along the way $\mathcal{S}_{0:L}$. Since causal states are sufficient statistics for prediction, the future symbol X_L depends only on the most recent causal state \mathcal{S}_L and, finally, the optimality of ϵ -machines means that the next symbol can be predicted at the entropy rate h_μ .

In other words, the state-block entropy that employs a process's ϵ -machine presentation is a straight line with slope h_μ and y -intercept $H[\mathcal{S}_0X_{0,0}] = H[\mathcal{S}_0] \equiv C_\mu$. Note that $H[\mathcal{S}_0X_{0:L}] \geq H[X_{0:L}]$ with equality if and only if $H[\mathcal{S}_0|X_{0:L}] = 0$. Since conditioning never increases uncertainty, these two block-entropy curves remain equal from that point onward. This necessarily implies that

they tend to the same asymptote. So, if the state-block entropy curve ever equals the block entropy curve, then the y -intercepts of each curve must also be equal: $C_\mu = \mathbf{E}$. Stated differently, the two curves meet if and only if the process has $\chi = 0$.

C. Block-state entropy

Finally, we consider the block-state entropy $H[X_{0:L}|\mathcal{R}_L]$, a measure of the joint uncertainty one has in a block of symbols and the presentation's subsequent internal state. Once again, our interest here is with ϵ -machines, and so we consider $H[X_{0:L}|\mathcal{S}_L]$. Unlike the state-block entropy, however, the behavior of this entropy is nontrivial. We recall a number of its properties and also establish the equivalence of the cryptic order definitions given in Refs. [9, 12]. Then, we provide a detailed proof of its convexity, as this does not appear previously.

The block-state entropy begins at C_μ when $L = 0$. As L increases, the curve is nondecreasing and tends, from above, to the same linear asymptote as the block entropy: $\mathbf{E} + h_\mu L$. Since the state-block entropy is $C_\mu + h_\mu L$ and since $C_\mu \geq \mathbf{E}$, we see that the state-block entropy curve is greater than or equal to the block-state entropy: $H[\mathcal{S}_0|X_{0:L}] \geq H[X_{0:L}|\mathcal{S}_L]$. Equality for $L > 0$ occurs if and only if the process has $C_\mu = \mathbf{E}$ or, equivalently, $\chi = 0$ and, then, the curves are equal for all L .

Similarly, the block-state entropy is greater than or equal to the block entropy: $H[X_{0:L}|\mathcal{S}_L] \geq H[X_{0:L}]$. We have equality if and only if $H[\mathcal{S}_L|X_{0:L}] = 0$. Recall, the smallest such L is the Markov order R . So, the block-state entropy equals the block entropy only at the Markov order. Further, once the curves are equal, they remain equal:

$$H[X_{0:L}|\mathcal{S}_L] = H[X_{0:L}] \Rightarrow H[X_{0:L+1}|\mathcal{S}_{L+1}] = H[X_{0:L+1}].$$

This can shown by individually expanding both $H[X_{0:L+1}|\mathcal{S}_{L+1}]$ and $H[X_{0:L+1}]$ to $H[X_{0:L}] + h_\mu$. The interpretation is that the two curves become equal only at the Markov order and only after both curves have reached their linear asymptotes.

Reference [12] defined the cryptic order as the minimum L for which the block-state entropy reaches its asymptote. This is in contrast to the definition provided here and also in Ref. [9], which defines the cryptic order as the minimum L for which $H[\mathcal{S}_L|X_{0:}] = 0$. We now establish the equivalence of these two definitions.

Theorem 1.

$$H[\mathcal{S}_L|X_{0:}] = 0 \iff H[X_{0:L}|\mathcal{S}_L] = \mathbf{E} + h_\mu L. \quad (1)$$

Proof.

$$H[\mathcal{S}_L|X_{0:}] = 0 \quad (2)$$

$$\iff H[\mathcal{S}_0|X_{0:}] = H[\mathcal{S}_0|X_{0:L}, \mathcal{S}_L] \quad (3)$$

$$\iff I[\mathcal{S}_0; X_{0:}] = I[\mathcal{S}_0; X_{0:L}, \mathcal{S}_L] \quad (4)$$

$$\iff \mathbf{E} = H[X_{0:L}, \mathcal{S}_L] - H[X_{0:L}, \mathcal{S}_L|\mathcal{S}_0] \quad (5)$$

$$\iff H[X_{0:L}, \mathcal{S}_L] = \mathbf{E} + H[\mathcal{S}_L|\mathcal{S}_0, X_{0:L}] + H[X_{0:L}|\mathcal{S}_0] \quad (6)$$

$$\iff H[X_{0:L}, \mathcal{S}_L] = \mathbf{E} + h_\mu L. \quad (7)$$

The step from Eq. (2) to Eq. (3) follows from Thm. 1 of Ref. [12]. In moving from Eq. (4) to Eq. (5), we used the prescience of causal states $\mathbf{E} = I[\mathcal{S}_0; X_{0:}]$ [20]. Finally, Eq. (6) leads to Eq. (7) using unifilarity of ϵ -machines ($H[\mathcal{S}_L|\mathcal{S}_0, X_{0:L}] = 0$) and that they allow for prediction at the process entropy rate: $H[X_{0:L}|\mathcal{S}_0] = h_\mu L$. \square

We obtain estimates for the crypticity χ by considering the difference between the state-block and block-entropy curves:

$$\chi(L) \equiv H[\mathcal{S}_0|X_{0:L}] - H[X_{0:L}|\mathcal{S}_L] \quad (8)$$

$$= h_\mu L - H[X_{0:L}|\mathcal{S}_L]. \quad (9)$$

Ref. [9] showed that this approximation limits from below in a nondecreasing manner to the process crypticity: $\chi(L) \rightarrow \chi$ and $\chi(L) \leq \chi(L+1) \leq \chi$. This also provides an upper-bound estimate of the excess entropy:

$$\mathbf{E} \leq C_\mu - \chi(L).$$

Combined with the lower-bound estimate the block entropy provides, one can be confident in the estimates of excess entropy.

The retrodictive error $H[X_{0:L}|\mathcal{S}_L]$ is the difference of the block-state entropy from the statistical complexity. It is also the difference of $\chi(L)$ from $h_\mu L$. Furthermore, it follows from Ref. [12] that the asymptotic retrodiction rate [24] is equal to the process entropy rate:

$$\lim_{L \rightarrow \infty} \frac{H[X_{0:L}|\mathcal{S}_L]}{L} = h_\mu.$$

In a sense, this describes short-term retrodiction. As we will see in a moment, order- R spin-chains are a class of processes that have no retrodiction error for a full R -block. The opposite class, in this sense, consists of processes with $\chi = 0$ —that is, the co-unifilar processes. These immediately begin retrodiction at the optimal rate, which is h_μ .

Finally, we establish the convexity of the block-state entropy, which appears to be new.

Theorem 2. $H[X_{0:L}|\mathcal{S}_L]$ is convex upwards in L .




FIG. 6. Four variable I-diagram for the block-state entropy convexity proof, with the needed sigma-algebra atoms appropriately labeled.

Proof. Convexity here means:

$$\begin{aligned} H[X_{0:L+1}S_{L+1}] - H[X_{0:L}S_L] \\ \geq H[X_{0:L}S_L] - H[X_{0:L-1}S_{L-1}] . \end{aligned}$$

Stationarity gives us:

$$\begin{aligned} H[X_{-1:L}S_L] - H[X_{0:L}S_L] \\ \geq H[X_{-1:L-1}S_{L-1}] - H[X_{0:L-1}S_{L-1}] . \end{aligned}$$

Simplifying, we have:

$$H[X_{-1}|X_{0:L}S_L] \geq H[X_{-1}|X_{0:L-1}S_{L-1}] .$$

We can use the I-diagram of Fig. 6 to help understand this last convexity statement. There, it translates into:

$$\alpha + \gamma \geq \alpha + \beta ,$$

or, since $\alpha \geq 0$:

$$\gamma \geq \beta . \quad (10)$$

Using the fact that the causal state is an optimal representation of the past, we have the following expressions that are asymptotically equivalent to the entropy rate h_μ :

$$\begin{aligned} H[X_{L-1}S_L|S_{L-1}] &= \beta + \epsilon + \delta + \zeta \\ H[X_{L-1}S_L|S_{L-1}X_{0:L-1}] &= \beta + \zeta \\ H[X_{L-1}S_L|S_{L-1}X_{-1}] &= \epsilon + \zeta \\ H[X_{L-1}S_L|S_{L-1}X_{-1}X_{0:L-1}] &= \zeta . \end{aligned}$$

The associations with the sigma-algebra atoms are readily gleaned from the I-diagram. Note that the various finite- L expressions for the entropy rate rely on the shielding property of the causal states and also on the ϵ -machine's unifilarity. Taken together in the $L \rightarrow \infty$ limit, the four

relations yield:

$$\begin{aligned} \zeta &= h_\mu \text{ and} \\ \beta &= \delta = \epsilon = 0 . \end{aligned}$$

These, in turn, transform the convexity criterion of Eq. (10) into the simple statement that:

$$\gamma \geq 0 .$$

Since $\gamma = I[X_{-1}; S_{L-1}|X_{0:L}S_L]$ is a conditional mutual information and, therefore, positive semidefinite, this establishes that the block-state entropy is convex. \square




FIG. 7. The block $H[X_{0:L}]$, state-block $H[S_0X_{0:L}]$, and block-state entropy $H[X_{0:L}S_0]$ curves compared. The sloped dashed line is the asymptote $\mathbf{E} + h_\mu L$, to which both the block entropy and state-block entropy asymptote. Finite Markov order and finite cryptic order are illustrated by the vertical dashed lines that indicate where the entropies meet the linear asymptote, respectively. The convergence of the crypticity approximation $\chi(L)$ to χ is also shown.

It will help to summarize the point that we have now reached. We used the various block entropy curves to synthesize much of our information-theoretic viewpoint of a process into a single representation—that shown in Fig. 7. We can amortize the effort to develop this viewpoint by applying it to a broad class of processes familiar from statistical mechanics.

VI. CRYPTICITY IN SPIN CHAINS

We first consider a subset of processes drawn from statistical mechanics known as one-dimensional *spin chains*. (For background, see Refs. [19, 25].) They are processes

such that $H[X_{0:R}] = C_\mu$. Using this, the simple geometry presented in Fig. 7 reveals that:

$$\chi(L) = \begin{cases} h_\mu L & 0 \leq L \leq R, \\ \chi & L > R, \end{cases} \quad (11)$$

and this, in turn, implies that $k = R$. This can also be seen through Eq. (9). Recall, the block-state entropy is nondecreasing and begins at C_μ . Since spin chains have $H[X_{0:R}] = C_\mu$, we know that the block-state entropy curve for spin chains must remain flat until $L = R$. Consequently, $H[X_{0:L}|\mathcal{S}_L] = 0$ and $\chi(L) = h_\mu L$ for $L \leq R$. Notice that $H[X_{0:R}|\mathcal{S}_R]$ not vanishing gives a way to understand how $\chi(L)$ deviates from linear growth. That is, the nonlinearity of the approach of $\chi(L)$ to χ is exactly the coentropy $H[X_{0:L}|\mathcal{S}_L]$.

This property is tantamount to a very simple test to determine if a process is a spin chain. If one obtains a plot similar to Fig. 7 for the process in question, it is a spin chain if $H[X_{0:L}, \mathcal{S}_L]$ goes from $(0, C_\mu)$ flat to (R, C_μ) , and then follows $\mathbf{E} + Lh_\mu$. That is, the block-state entropy curve is flat until it reaches its asymptote at $L = k = R$, at which point it tracks it.

Furthermore, given (i) the above proof, (ii) the concavity proof from Sec. V C, and (iii) the fact that $k \leq R$, for a given \mathbf{E} , h_μ , and R spin chains are seen to be maximally-cryptic processes. By this we mean that for all processes with a particular set of values for \mathbf{E} , h_μ , and R , the process that maximizes χ is a spin chain. This implies that C_μ is also maximized.




FIG. 8. An order-2 Markov spin chain with full support.

Figures 8 and 9 show two order-2 Markov spin chains. The first is a full-support order-2 Markov chain, while the second has only partial support. In fact, the latter process has the Golden Mean support consisting of all bi-infinite sequences that do not contain consecutive 0s.

Figure 10 gives an ϵ -machine of similar structure to the spin chains just examined and, while it is also an order-2 Markov process, it is not a spin chain. The reason is that one causal state (labeled “01, 11”) is induced by two




FIG. 9. An order-2 Markov spin chain with partial support.




FIG. 10. An order-2 Markov process, but not a spin chain.

words: 01 and 11. This means that the correspondence between inducing-words and causal states is broken. It is no longer a spin chain.

We close this section with a number of open questions about spin chains. The first two regard the structure of spin chains. If an ϵ -machine is a subgraph of an order- R Markov skeleton, then is it a spin chain? That is, does the removal of an edge from a spin chain produce another spin chain? The intuition behind this question is straightforward: Removing transitions disallows blocks, but it would not cause any block to be associated with a different state. A related question asks if all spin chains are of this form.

The next two questions regard the transformation from a spin chain to any other process and vice versa. First, can any order- R Markov, order- k cryptic ϵ -machine be obtained by starting with an order- R Markov skeleton, reducing some probabilities to zero and adjusting others to cause state merging? Also, given an order- R Markov, order- k cryptic ϵ -machine, we can break the existing degeneracy so that $H[X_{0:R}] = C_\mu$. How does the nonspin chain we started with compare with the spin chain we end up with?

VII. GEOMETRIC CONSTRAINTS

The geometry of the block entropy convergence illustrated in Fig. 7 can be exploited. In particular, as we will now show, a variety of constraints leads to further results on the allowed convergence behaviors the block and block-state entropy curves can express. Figure 11 depicts these results graphically.




FIG. 11. Constraints on entropy convergence, illustrated for a process that is order-5 Markov and order-4 cryptic. The blue region circumscribes where the block entropy curve can lie; the tan, where the block-state entropy may be. These and the discreteness of L lead to restrictions on allowed cryptic orders as well.

First, given the block entropy’s concavity and its asymptote, one sees that the block entropy curve is contained within the triangle described by $\{(0, 0), (0, \mathbf{E}), (R, H[X_{0:R}])\}$. We also know that the block entropy cannot grow faster than $H[X_0]$ and this excludes the triangle $\{(0, 0), (0, \mathbf{E}), (1, H[X_0])\}$. The resulting allowed region is shown in light blue in Fig. 11.

Second and similarly, the block-state entropy’s own properties require it to be within a triangle described by $\{(0, C_\mu), (\frac{\chi}{h_\mu}, C_\mu), (R, H[X_{0:R}])\}$.

Third, since the entropy functions are defined for discrete values of word length L , we can go a little further than these observations. The block-state entropy cannot intersect the asymptote $\mathbf{E} + h_\mu L$ at a noninteger L . Therefore the small triangle $\{(\lfloor \frac{\chi}{h_\mu} \rfloor, C_\mu), (\frac{\chi}{h_\mu}, C_\mu), (\lceil \frac{\chi}{h_\mu} \rceil, \mathbf{E} + \lceil \frac{\chi}{h_\mu} \rceil h_\mu)\}$ is excluded. The resulting allowed trapezoid is displayed in tan in Fig. 11.

Fourth, recalling results on the block-state entropy, this exclusion means that processes with $C_\mu \neq \mathbf{E} + h_\mu k$, for some k , must have a degree of nonoptimal retrodiction. In short, they are prevented from being spin chains.

Finally, given a process that has cryptic order k , we see that $C_\mu \leq \mathbf{E} + h_\mu k$. A more detailed result then says that $C_\mu = \mathbf{E} + h_\mu k$ if and only if $H[X^k] = C_\mu$. Moreover, it is Markov order- k ; that is, it is a spin chain.

VIII. THE CRYPTIC MARKOVIAN ZOO

It turns out that there exist finite-state processes with all combinations of Markov and cryptic order; subject, of course, to the constraint that $R \geq k$. These range from the zero structural complexity independent, identically distributed processes, for which $R = 0$ and $k = 0$, to few-state processes where either or both are infinite. (For a complementary and exhaustive survey see Ref. [26].) In practice, given what we now know about these properties, it is not difficult to design a variety of processes that fulfill a given specification.

Also noteworthy is how the introduction of the new crypticity “coordinate” affects our view of several well studied examples. For instance, the Even Process [6] is one of the canonical finite-state, infinite-order Markov processes. In the past, it was often thought of as representing both intractability and compactness. Now, though, we see that it is trivial, being 0-cryptic. The Golden Mean Process, one of the simplest (order-1 Markov) subshifts of finite-type studied is now seen as more sophisticated, being 1-cryptic. These and similar explorations naturally lead one to delve deeper to find extreme examples—such as the Nemo process below—that are infinite in both cryptic and Markov orders. Again, see Ref. [26].

Figure 12 presents a crypticity-Markovity roadmap for the space of finite-state processes. Borrowing from the immediately preceding citations, it also displays a select few processes using their ϵ -machines to show concretely the full diversity of possible Markov and cryptic orders a finite-state process can possess. The green bar at $k = 0$ consists of all co-unifilar processes. The orange line contains all processes where the Markov and cryptic orders are identical—a subset of which are the spin chains. All other processes lie above this line. The Even Process is in the upper left corner. The Golden Mean Process (no consecutive 0s) is in the lower left. The ∞ -cryptic, infinite-order Markov Nemo Process is in the upper right corner. Several of the other prototype ϵ -machines depicted illustrate (R, k) -parametrized classes of process for whom the Markov and cryptic orders can be selected arbitrarily.




FIG. 12. The crypticity-Markovity roadmap for finite-state stationary processes: The range of possible Markov and cryptic orders, illustrated by a sample of processes depicted by their ϵ -machines. Lower left: The Fair Coin Process and all other IID processes. Upper left: The ∞ -cryptic Even Process. Upper right: The Nemo Process. Left vertical (green) line: The co-unifilar processes.

IX. INFORMATION DIAGRAMS FOR STATIONARY PROCESSES

Information diagrams, or simply *I-diagrams*, are an important analysis tool in using information theory to analyze multivariate stochastic processes [27]. They are particularly useful when working with processes and, as we have already seen here, give a good deal of insight when the ϵ -machine presentation is employed [8, 11].

The essential idea is that there is a one-to-one correspondence between information-theoretic quantities—mutual information and conditional and joint entropies—and measurable sets. Constructively, informational relationships and constraints are depicted via set-theoretic operations: joint entropies are set unions, conditional entropies correspond to set difference, mutual information corresponds to set intersection, and the like. The math-

ematical structure is a sigma algebra over the process's events (words). The noncomposite sets are the *atoms* of the sigma algebra and their size is the magnitude of the corresponding informational quantities. When depicted graphically, though, one often ignores magnitudes and, instead, focuses on the set-theoretic relationships.

Armed with simple and familiar rules, one can often accomplish several algebraic calculational steps on compound entropy expressions via a simple I-diagram and a small description. Perhaps more importantly, I-diagrams afford a visual calculus that lends a heightened intuition about complicated relationships among random variables.

Figures 13 through 17 show how to make more explicit and intuitive the preceding formal views of entropy convergence and its relationship to Markovity and crypticity. The two large circles in each represent the past via

$H[X_{:0}]$ and the future via $H[X_{0:}]$. The excess entropy $\mathbf{E} = I[X_{:0}; X_{0:}]$, being a mutual information, is their intersection. The I-diagrams there show the nested dependence of the various information measures as one increases block size and so increases the number of random variables. In the general multivariate case this would lead to an explosion of atoms. However, due to the nature of processes and the ϵ -machine itself many simplifications are possible. Figures 14-16 also depict the ϵ -machine's causal-state information, $C_\mu = H[\mathcal{S}]$, as a circle entirely inside the past $H[X_{:0}]$. This is so, since the causal states are a function of the past.




FIG. 13. Information diagram for an order-4 Markov process. Only the four most recent history symbols are needed to reduce as much uncertainty in the future as using the whole past would.

To start with the simplest case, Fig. 13 gives the I-diagram for an order-4 Markov process. As one expects, only the four most recent history symbols are needed to reduce as much uncertainty in the future as using the whole past would. Equivalently, as soon as the history contains four symbols, all of the shared information between the past and the future (the excess entropy \mathbf{E}) is captured.




FIG. 14. Causal state is overlaid onto an I-diagram for an order-4 Markov process. As drawn, no fewer than 4 history symbols are required to determine the causal state. The causal state, though, does not generally determine this length-four history.

Figure 14 then overlays the causal state measure $H[\mathcal{S}]$.

In this, we see that no fewer than four history symbols are required to determine the causal state. Importantly, it is now also made explicit that causal states do not generally determine this history.




FIG. 15. An I-diagram for an order-4 Markov process, but order-3 cryptic. Four history symbols are required to determine the state, but only three are required if one conditions on the future.

Consider now the order-4 Markov, order-3 cryptic process of Figure 15. As before, four history symbols are required to determine the state. But, as depicted, only three history symbols are required if one conditions on the future as well.




FIG. 16. The separation between Markov and cryptic orders can be widened: A Markov order-4, cryptic order-2 process.

Figure 16 demonstrates how the difference between Markov and cryptic orders can be increased without bound. The I-diagram illustrates the sigma-algebra for an order-4 Markov, order-2 cryptic process.

Finally, Fig. 17 gives the I-diagram for an order-4 spin chain. Several features of spin chains are clearly rendered in this I-diagram. First, the shortest history that uniquely determines the state occurs at length 4. Specifically, as depicted, $\min_L : H[\mathcal{S}_0 | X_{-L:0}] = 4$. And, at the same time, this length-4 history is itself uniquely determined by the causal state.




FIG. 17. The highly regular I-diagram for an order-4 spin chain.

X. CONCLUSION

Crypticity, as the difference between a process's stored information and its observed information, is a key property. The fundamental definitions, Eqs. (1) and (2), though, are not immediately transparent. However, they do lead to several interpretations that prove useful in different settings. Given this, our main goals were to explicate the basic notions behind crypticity and to motivate various of its interpretations. Along the way, we provided a new geometric interpretation for cryptic order, established a number of previously outstanding properties, and illustrated crypticity by giving a complete analysis for spin chains.

More specifically, using state-paths, we introduced several new interpretations of crypticity that not only helped to explain the basic idea but also suggest future applications in distributed dynamical systems. We also gave a simple geometric picture that relates cryptic and Markov orders. We established the equivalence between co-unifilarity and being 0-cryptic, as well as the concavity of the block-state entropy $H[X_{0:L}\mathcal{S}_L]$. We derived several geometric constraints and drew out their implications for bounds on crypticity. These also led to an improved bound on Markov order. Presumably, the bounds will help improve estimates of crypticity and cryptic order, in both the finite and infinite cases.

To give a sense of the relationship between cryptic and Markov orders we gave a graphical overview classifying processes in their terms. In a complementary way, we introduced the technique of foliated information diagrams to analyze entropy convergence and Markov and cryptic orders in terms of Shannon information measures and their now block-length-dependent sigma algebra.

To ground the results in a concrete and familiar class of processes we analyzed range- R 1D spin chains in detail. We established their Markov order and showed that the block-state entropy $H[X_{0:L}\mathcal{S}_L]$ is flat for spin chains

and that $\chi(L) = Lh_\mu$, for all $L \leq R$. From these properties one can determine whether or not a given process is representable as a spin chain: Is the R -block entropy equal to the statistical complexity? The properties also suggest what the processes in the neighborhood of a spin chain look like.

Finally, by way of making contact with applications to physics and computation, we close by briefly outlining the relationship between crypticity and dynamical irreversibility in physical processes [28]. Consider the *morph map* $\phi : \mathcal{S}_0 \rightarrow \{X_{0:}\}$. A process's entropy rate controls the prediction uncertainty of this map: $h_\mu \equiv \lim_{L \rightarrow \infty} H[X_{0:L}|\mathcal{S}_0]$. Now, consider the state uncertainty determined by the inverse of the morph map: $\phi^{-1} : X_{0:} \rightarrow \{\mathcal{S}_0\}$. This is already familiar. The crypticity controls this uncertainty: $\chi = \lim_{L \rightarrow \infty} H[\mathcal{S}_0|X_{0:L}]$. Just as the entropy rate is a process's rate of producing information, the crypticity is its rate of information loss or, what one can call, a process's *information-processing irreversibility*. And the latter, appropriately adapting Landauer's Principle [29], provides a lower bound on the energy dissipation required to support a process's irreversible intrinsic computation. We leave the full development of the thermodynamics of intrinsic computation, however, to another venue.

ACKNOWLEDGMENTS

This work was partially supported by NSF Grant No. PHY-0748828 and by the Defense Advanced Research Projects Agency (DARPA) Physical Intelligence Subcontract No. 9060-000709. The views, opinions, and findings contained in this article are those of the authors and should not be interpreted as representing the official views or policies, either expressed or implied, of the DARPA or the Department of Defense.

Appendix A: Why Crypticity?

There are many ways to assemble information-theoretic quantities—more specifically, information measures [27]. Why should one care about crypticity and cryptic order? What makes them special? We show that crypticity stands out among reasonable alternative measures by a rather direct comparison.

It turns out that there are fewer information quantities than one might expect—at least fewer interesting ones—over pasts, futures, and states. Let's limit ourselves to quantities that depend on only a finite set of objects and require that we look for a “1-parameter finitization” property, based on block length. In this case, we

can make an exhaustive list of the information measures and describe each one. The list, at first, appears long. But this length is illustrative of the fact that crypticity and cryptic order really do capture a relatively unique process property. Everything else is either trivial, periodic, or Markov.

Table I presents the list. It was assembled in a direct way by systematically writing down alternative expressions over single variables, pairs of variables and their joint and conditional entropy possibilities, over three variables, and so on. One could also consider enumerating only the relevant sigma-algebra atoms. This, however, obscures parallels to existing quantities.

In addition, alternatives such as $H(X_{-L:0}|X_{:0})$ are not included, since they are trivial. Nor were quantities such as $H(X_{:0}|X_{-L:0})$ added, although they could be. Quantities along these lines would needlessly expand the list, to little benefit.

As elsewhere here, we assume the state random variable denotes a causal state.

Appendix B: Equivalence of Forward and Reverse Restricted State-Paths

Why are the restricted state-paths the same in the forward and backward lattice diagrams of Figs. 3 and 4? Recall that a forward path is allowed if $\Pr(X_{0:L} = w, \mathcal{S}_1 = \sigma_B | \mathcal{S}_0 = \sigma_A) \neq 0$. Similarly, a backward path is allowed when $\Pr(\mathcal{S}_0 = \sigma_A, X_{0:L} = w | \mathcal{S}_L = \sigma_B) \neq 0$. Since both causal states σ_A and σ_B have nonzero probability by definition of being recurrent, we see that we can state both cases as paths for which $\Pr(\mathcal{S}_0 = \sigma_A, X_{0:L} = w, \mathcal{S}_L = \sigma_B) \neq 0$.




FIG. 18. Why forward and backward restricted paths are the same. In this figure state-paths are traced back from final states. Cf. Fig. 4.

Figure 18 illustrates this by tracing state-paths backward through the machine starting at each final state. Of course, since processes and their ϵ -machines are generally not counifilar, there will be splitting in these paths. For example, consider the paths that end in state A on a 1. A's predecessors on a 1 are states A and E.

Note that this produces a different initial set of candidate state-paths, when compared with those in light blue in Fig. 4. Now, eliminate all paths that do not trace back

Information Measure	Property Detected
$H[A]$	
$H[X_{:0}] = H[X_{-L:0}]$	Periodic
$H[X_{0:}] = H[X_{0:L}]$	Periodic
$H[A B]$	
$H[\mathcal{S} X_{:0}] = H[\mathcal{S} X_{-L:0}]$	Markov
$H[\mathcal{S} X_{0:}] = H[\mathcal{S} X_{0:L}]$	Markov
$H[X_{0:} X_{:0}] = H[X_{0:} X_{-L:0}]$	Markov
$H[X_{:0} X_{0:}] = H[X_{:0} X_{0:L}]$	Markov
$H[X_{:0} \mathcal{S}] = H[X_{-L:0} \mathcal{S}]$	Periodic
$H[X_{0:} \mathcal{S}] = H[X_{0:L} \mathcal{S}]$	Periodic
$H[A BC]$	
$H[\mathcal{S} X_{:0}X_{0:}] = H[\mathcal{S} X_{-L:0}X_{0:}]$	Cryptic Order
$H[\mathcal{S} X_{:0}X_{0:}] = H[\mathcal{S} X_{:0}X_{0:L}]$	Trivial
$H[X_{:0} \mathcal{S}X_{0:}] = H[X_{:0} \mathcal{S}X_{0:L}]$	Trivial
$H[X_{0:} X_{:0}\mathcal{S}] = H[X_{0:} X_{-L:0}\mathcal{S}]$	Trivial
$H[X_{:0} \mathcal{S}X_{0:}] = H[X_{-L:0} \mathcal{S}X_{0:}]$	Periodic
$H[X_{0:} X_{:0}\mathcal{S}] = H[X_{0:L} X_{:0}\mathcal{S}]$	Periodic
$H[AB]$	
$H[X_{:0}\mathcal{S}] = H[X_{-L:0}\mathcal{S}]$	Periodic
$H[\mathcal{S}X_{0:}] = H[\mathcal{S}X_{0:L}]$	Periodic
$H[X_{:0}X_{0:}] = H[X_{-L:0}X_{0:}]$	Periodic
$H[X_{:0}X_{0:}] = H[X_{:0}X_{0:L}]$	Periodic
$H[AB C]$	
$H[X_{:0}\mathcal{S} X_{0:}] = H[X_{-L:0}\mathcal{S} X_{0:}]$	Periodic
$H[\mathcal{S}X_{0:} X_{:0}] = H[\mathcal{S}X_{0:L} X_{:0}]$	Periodic
$H[X_{:0}\mathcal{S} X_{0:}] = H[X_{:0}\mathcal{S} X_{0:L}]$	Markov
$H[\mathcal{S}X_{0:} X_{:0}] = H[\mathcal{S}X_{0:} X_{-L:0}]$	Markov
$H[ABC]$	
$H[X_{:0}\mathcal{S}X_{0:}] = H[X_{-L:0}\mathcal{S}X_{0:}]$	Periodic
$H[X_{:0}\mathcal{S}X_{0:}] = H[X_{:0}\mathcal{S}X_{0:L}]$	Periodic

TABLE I. Alternative information measures over the past, the future, and the causal state, when they achieve their limit at finite block length L . As seen, almost all are either trivial, periodic, or detect the Markov property. Cryptic order stands out as unique.

successfully along the entire word. Fig. 18 shows these remaining state-paths in red and we see that they are the same as those in Fig. 4.

Appendix C: Crypticity and Co-unifilarity

Here, we explore the equivalence of $\mathbf{E} = C_\mu$, co-unifilarity, and 0-crypticity using several results obtained in Ref. [11]. With a small modification, the latter results allow for a more straightforward proof that leads to a better understanding of these relations.

The “forward” argument is that $\chi(L) = 0$ implies crypticity vanishes at all L . First, we recall two results.

Corollary 6 [11]: If there exists a $k \geq 1$ for which $\chi(k) = 0$, then $\chi(j) = 0$ for all $j \geq 1$.

Proposition 3 [11]: $\lim_{k \rightarrow \infty} \chi(k) = \chi$.

Combining Cor. 6 and Prop. 3, we have the following: If there exists a $k \geq 1$ for which $\chi(k) = 0$, then $\chi(j) = 0$ for all $j \geq 1$ and $\chi = 0$.

The “backward” argument is that vanishing in the

limit implies crypticity vanishes at all L .

Since $\chi(k)$ is nonnegative (conditional entropy) and nondecreasing (Prop. 2 [11]) and limits to χ (Prop. 3 [11]), we have that $\chi = 0$ implies $\chi(k) = 0$, for all $k \geq 0$.

All that remains is to recall that co-unifilarity is identical to $\chi(1) = 0$ and this establishes the desired chain of implications:

$$\begin{aligned} \text{Co-unifilar} &\iff \chi(1) = 0 \\ &\iff \exists k \geq 1 : \chi(k) = 0 \\ &\iff \chi(k) = 0, \forall k \geq 0 \\ &\iff \chi = 0 \\ &\iff \text{0-cryptic} . \end{aligned}$$

The heart of the result falls in the middle. It shows us that any nontrivial zero in $\chi(k)$ is equivalent to the entire function, as well as χ itself, vanishing.

-
- [1] J. P. Crutchfield, W. Ditto, and S. Sinha. Intrinsic and designed computation: Information processing in dynamical systems—Beyond the digital hegemony. *CHAOS*, 20(3):037101, 2010.
 - [2] G. E. Moore. Cramming more components onto integrated circuits. *Electronics*, 38(8):56–59, 1965.
 - [3] G. E. Moore. Progress in digital integrated electronics. *Technical Digest 1975. International Electron Devices Meeting, IEEE*, pages 11–13, 1975.
 - [4] G. E. Moore. Lithography and the future of moore’s law. *Proceedings of the SPIE*, 2437:1–8, May 1995.
 - [5] H. A. Atmanspracher and H. Scheingraber. *Information Dynamics*. Plenum, New York, 1991.
 - [6] J. P. Crutchfield and D. P. Feldman. Regularities unseen, randomness observed: Levels of entropy convergence. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 13(1):25–54, 2003.
 - [7] J. P. Crutchfield and K. Young. Inferring statistical complexity. *Phys. Rev. Lett.*, 63:105–108, 1989.
 - [8] J. P. Crutchfield, C. J. Ellison, and J. R. Mahoney. Time’s barbed arrow: Irreversibility, crypticity, and stored information. *Phys. Rev. Lett.*, 103(9):094101, 2009.
 - [9] J. R. Mahoney, C. J. Ellison, and J. P. Crutchfield. Information accessibility and cryptic processes. *J. Phys. A: Math. Theo.*, 42:362002, 2009.
 - [10] R. G. James, C. J. Ellison, and J. P. Crutchfield. Anatomy of a bit: Information in a time series observation. page submitted, 2010. Santa Fe Institute Working Paper 11-05-XXX; arxiv.org:1105.2988 [math.IT].
 - [11] C. J. Ellison, J. R. Mahoney, and J. P. Crutchfield. Prediction, retrodiction, and the amount of information stored in the present. *J. Stat. Phys.*, 136(6):1005–1034, 2009.
 - [12] J. P. Crutchfield, C. J. Ellison, J. R. Mahoney, and R. G. James. Synchronization and control in intrinsic and designed computation: An information-theoretic analysis of competing models of stochastic computation. *CHAOS*, 20(3):037105, 2010.
 - [13] J. P. Crutchfield and D. P. Feldman. Synchronizing to the environment: Information theoretic limits on agent learning. *Adv. in Complex Systems*, 4(2):251–264, 2001.
 - [14] A *synchronizing word* is symbol sequence that induces one and only one causal state. This is in contrast with a *minimal* synchronizing word, of which no proper prefix is a synchronizing word.
 - [15] The fact that any state is led by any symbol to at most one next state is the property known as unifilarity—a direct consequence of the states being causal states.
 - [16] The maze example is not a stationary process, so there are some important differences. For instance, there can be words that terminate *after* the length of the longest maze solution.
 - [17] Assume an N by N maze. A nonintersecting solution cannot contain more instructions than there are locations within the maze.
 - [18] To be more precise, it is not so much that the statistical complexity is derived from considering all paths as it is derived from considering *no* paths.
 - [19] J. P. Crutchfield and D. P. Feldman. Statistical complexity of simple one-dimensional spin systems. *Phys. Rev. E*, 55(2):1239R–1243R, 1997.
 - [20] C. R. Shalizi and J. P. Crutchfield. Computational mechanics: Pattern and prediction, structure and simplicity. *J. Stat. Phys.*, 104:817–879, 2001.
 - [21] J. R. Mahoney, C. J. Ellison, and J. P. Crutchfield. Information accessibility and cryptic processes: Linear combinations of causal states. 2009. arxiv.org:0906.5099 [cond-mat].
 - [22] N. Travers and J. P. Crutchfield. Asymptotically syn-

- chronizing to finitary sources. 2010. SFI Working Paper 10-06-XXX; arxiv.org:10XX.XXXX [XXXX].
- [23] One might choose to consider the process to be a bi-infinite collection of symbols or consider including causal states as well. We could similarly consider reverse causal states.
- [24] J. P. Crutchfield and C. R. Shalizi. Thermodynamic depth of causal states: Objective complexity via minimal representations. *Phys. Rev. E*, 59(1):275–283, 1999.
- [25] D. P. Feldman and J. P. Crutchfield. Discovering non-critical organization: Statistical mechanical, information theoretic, and computational views of patterns in simple one-dimensional spin systems. 1998. Santa Fe Institute Working Paper 98-04-026.
- [26] Ryan G. James, J. R. Mahoney, C. J. Ellison, and J. P. Crutchfield. Many roads to synchrony: Natural time scales and their algorithms. *submitted*, 2010. arxiv.org:1010.5545 [nlin.CD].
- [27] R. W. Yeung. A new outlook on Shannon’s information measures. *IEEE Trans. Info. Th.*, 37(3):466–474, 1991.
- [28] C. J. Ellison, J. R. Mahoney, R. G. James, J. P. Crutchfield, and J. Reichardt. Information symmetries in irreversible processes. pages 1–32, 2011. Santa Fe Institute Working Paper 11-07-028, arxiv.org:1107.2168 [cond-mat.stat-mech].
- [29] R. Landauer. Dissipation and noise immunity in computation, measurement, and communication. *J. Stat. Phys.*, 54(5/6):1509–1517, 1989.